

Heaps, Indexes and Execution Plans



Chennai SQL Server User Group

Blog: <http://sql-articles.com/category/cssug/>

MailDL : cssug@googlegroups.com

FB : www.facebook.com/groups/cssug/

Ramkumar Gopal

Living For SQL Server

Blog:

<http://www.sqlservercentral.com/blogs/livingforsqlserver/>



<http://www.facebook.com/LivingForSqlServer>

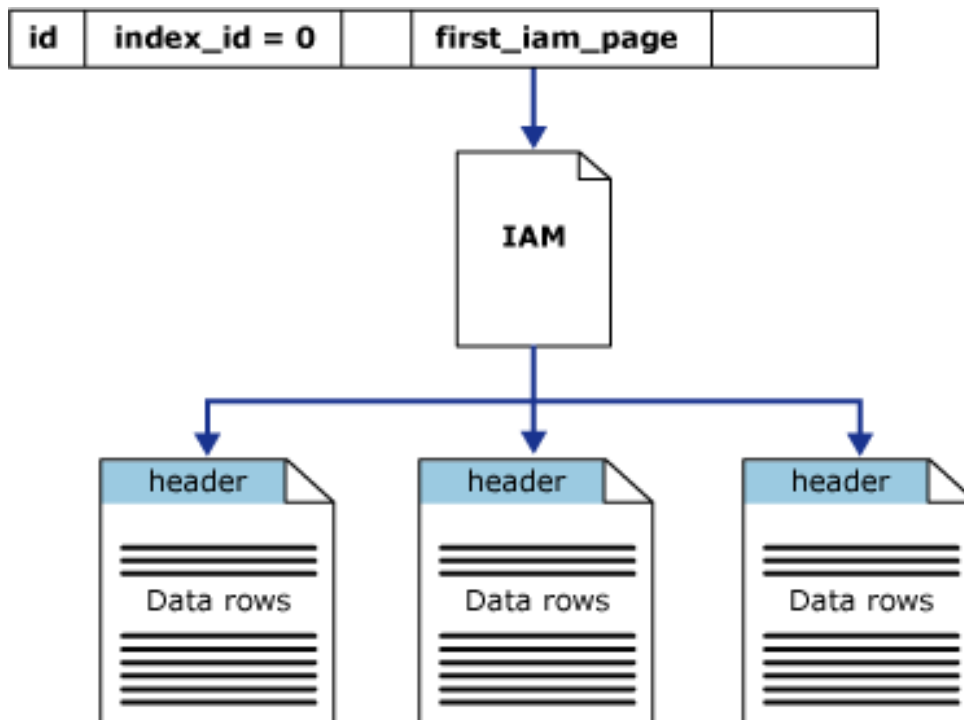


Agenda

- Heaps
- B-Trees
- Clustered Indexes
- Non Clustered Indexes
- Execution Plans

What is a Heap?

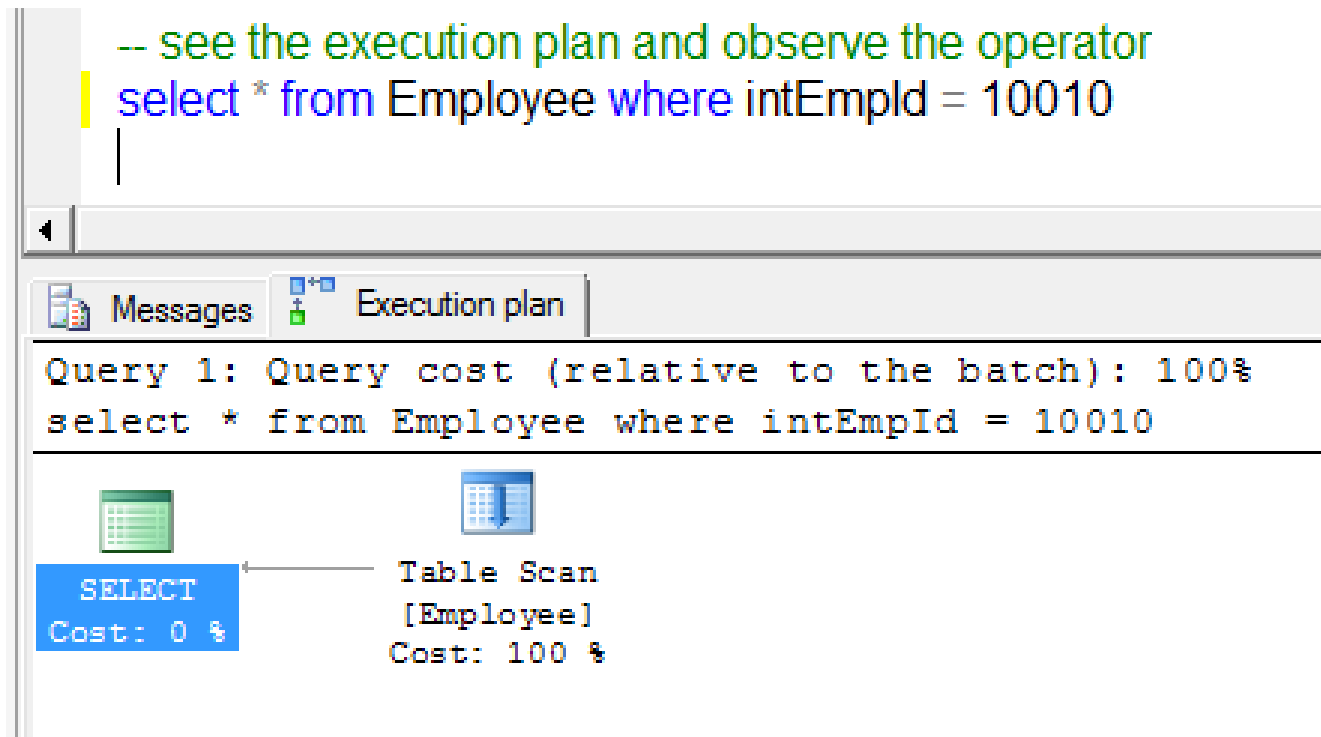
- * A heap is a table without a clustered index.
- * `index_id = 0`
- * SQL Server uses the IAM pages to move through the heap.
- * The data pages and the rows within them are not in any specific order and are not linked
- * The only logical connection between data pages is the information recorded in the IAM pages



Reading records from Heap

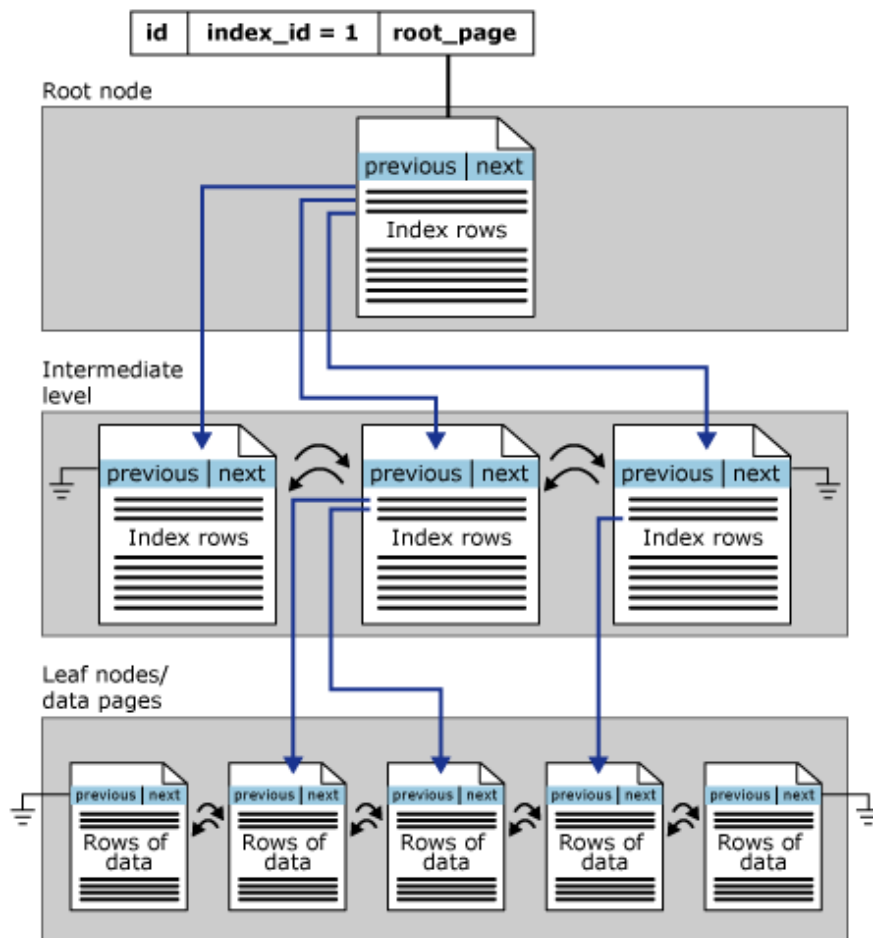
-- see the execution plan and observe the operator

```
select * from Employee where intEmpId = 10010
```



What is a B-Tree?

B-Tree – This structure helps to find the records quickly



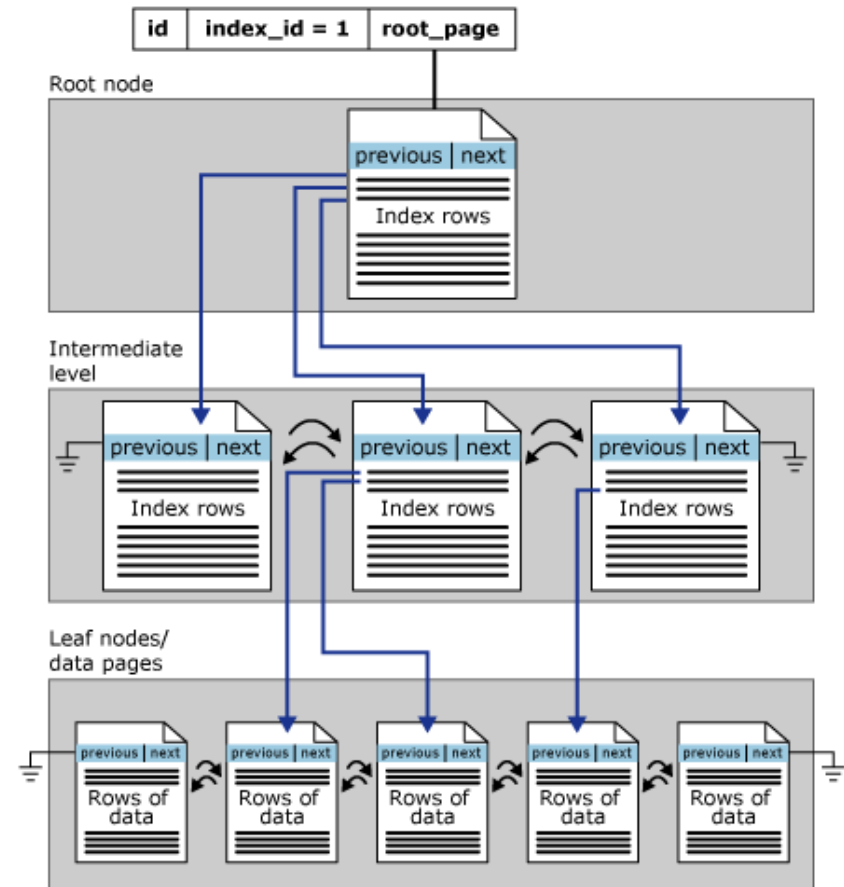
In SQL Server

Clustered Index and
Non Clustered Indexes

Are using B-Tree Structure

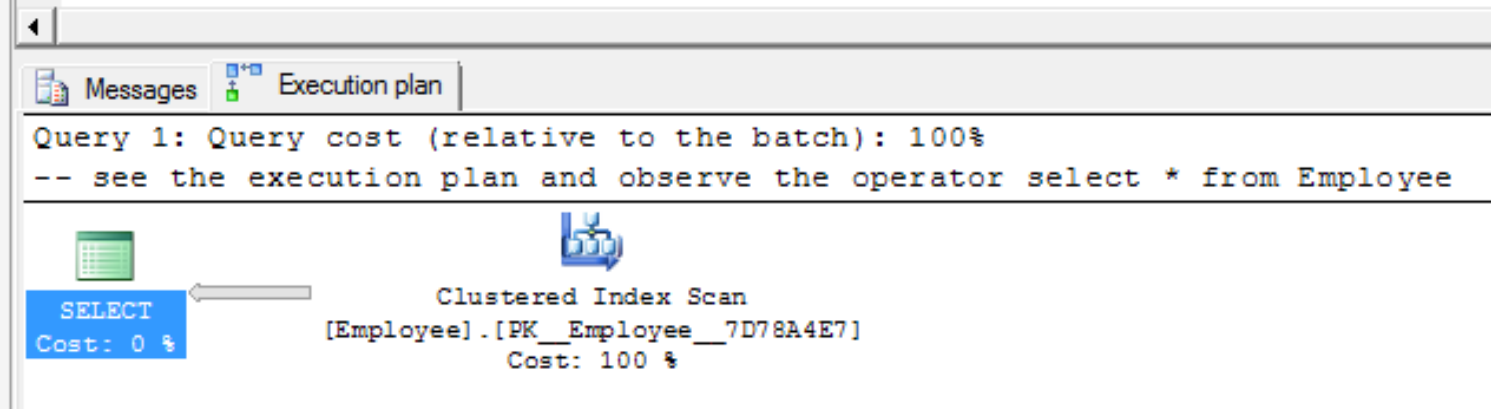
What is a Clustered Index?

- Index_id = 1
- Records are sorted based on CI Key
- Only one Clustered Index per table
- Root and Intermediate pages have Index keys
- Leaf node have actual data pages
- Carefully choose the CI Key

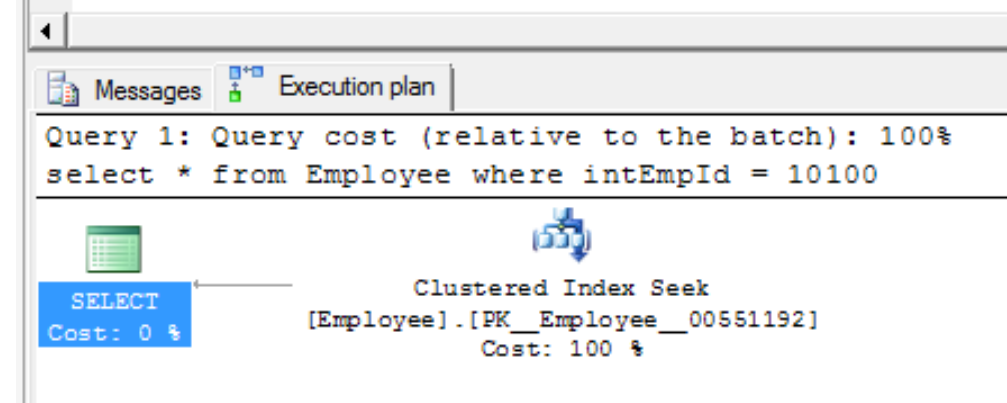


Reading records from Clustered Index

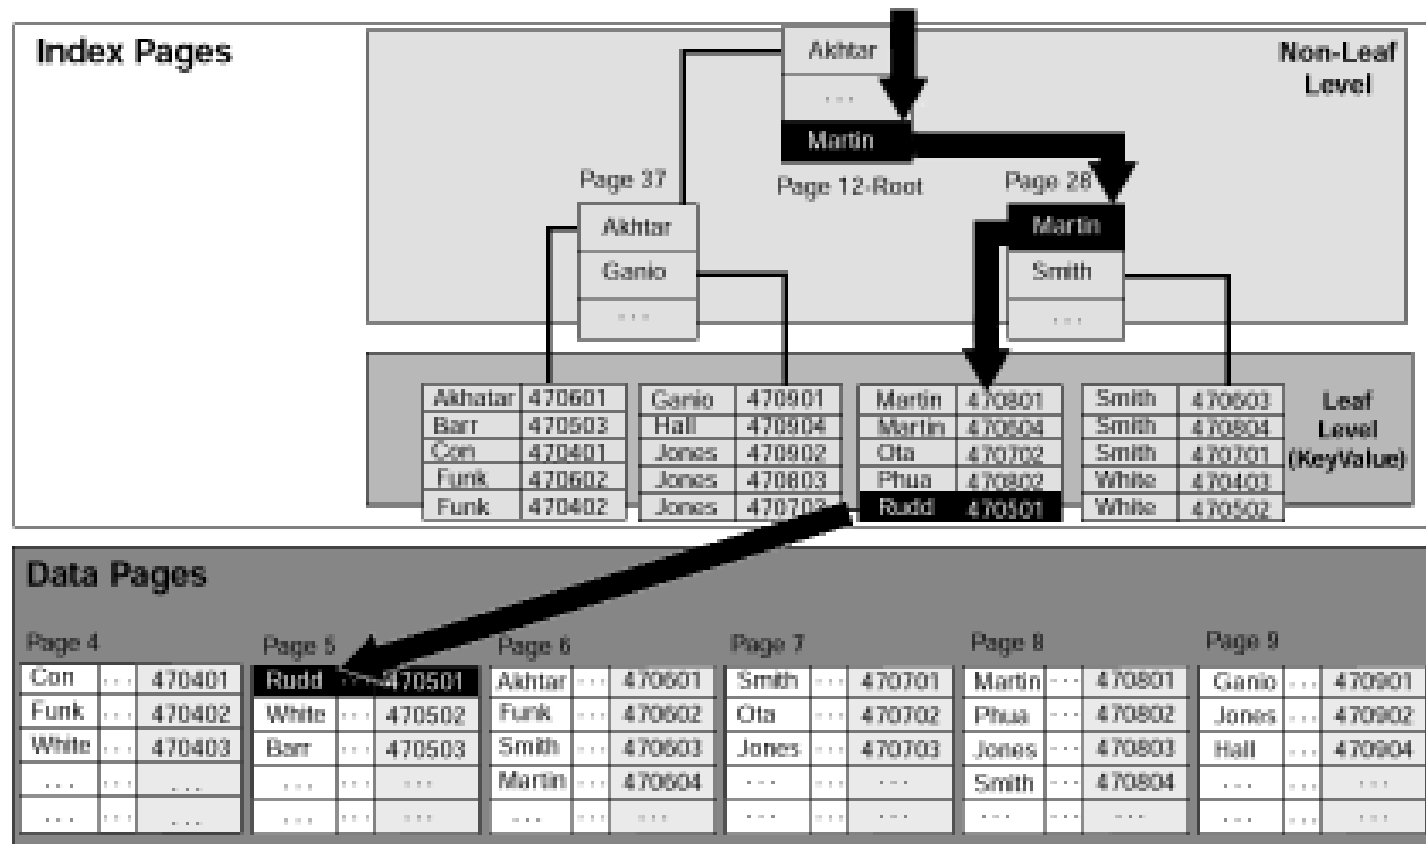
-- see the execution plan and observe the operator
`select * from Employee`



-- see the execution plan and observe the operator
`select * from Employee where intEmpId = 10100`

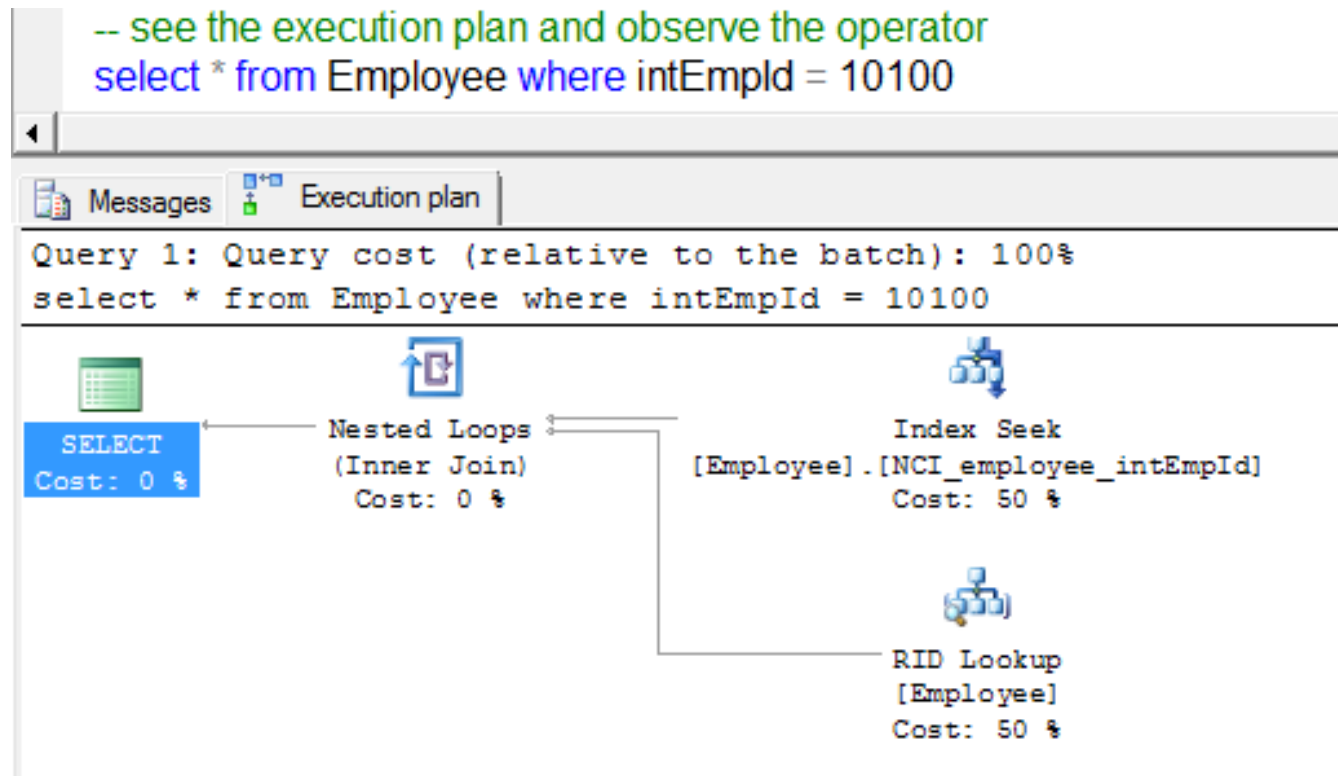


- Index_id > 1 and < 249
- Only Index entries in all B-Tree Levels
- Need to refer Heap or Clustered Index to fetch a record (Lookup)



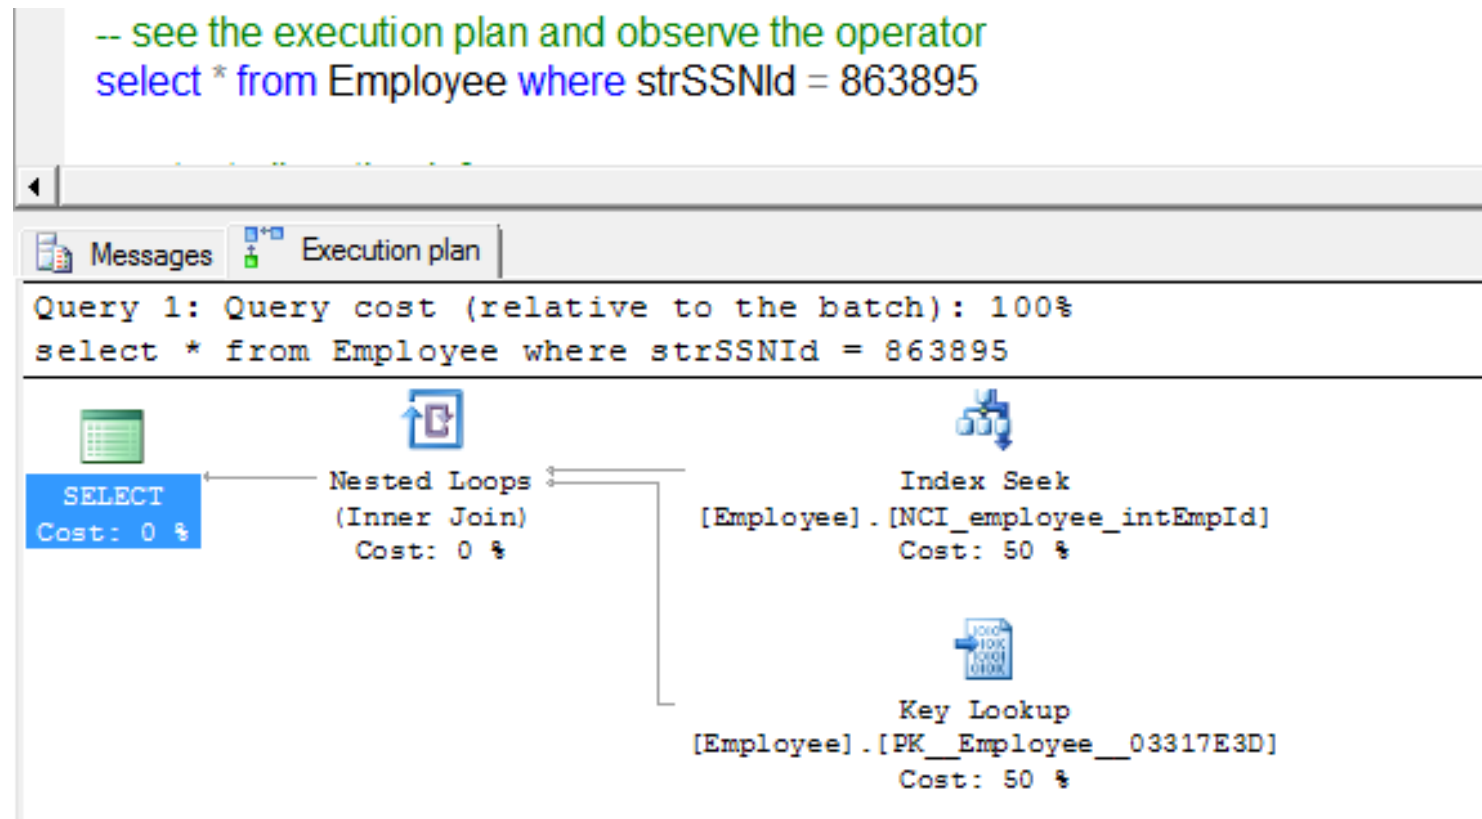
Reading records using Non Clustered Index

Reffering Heap with Non Clustered Index



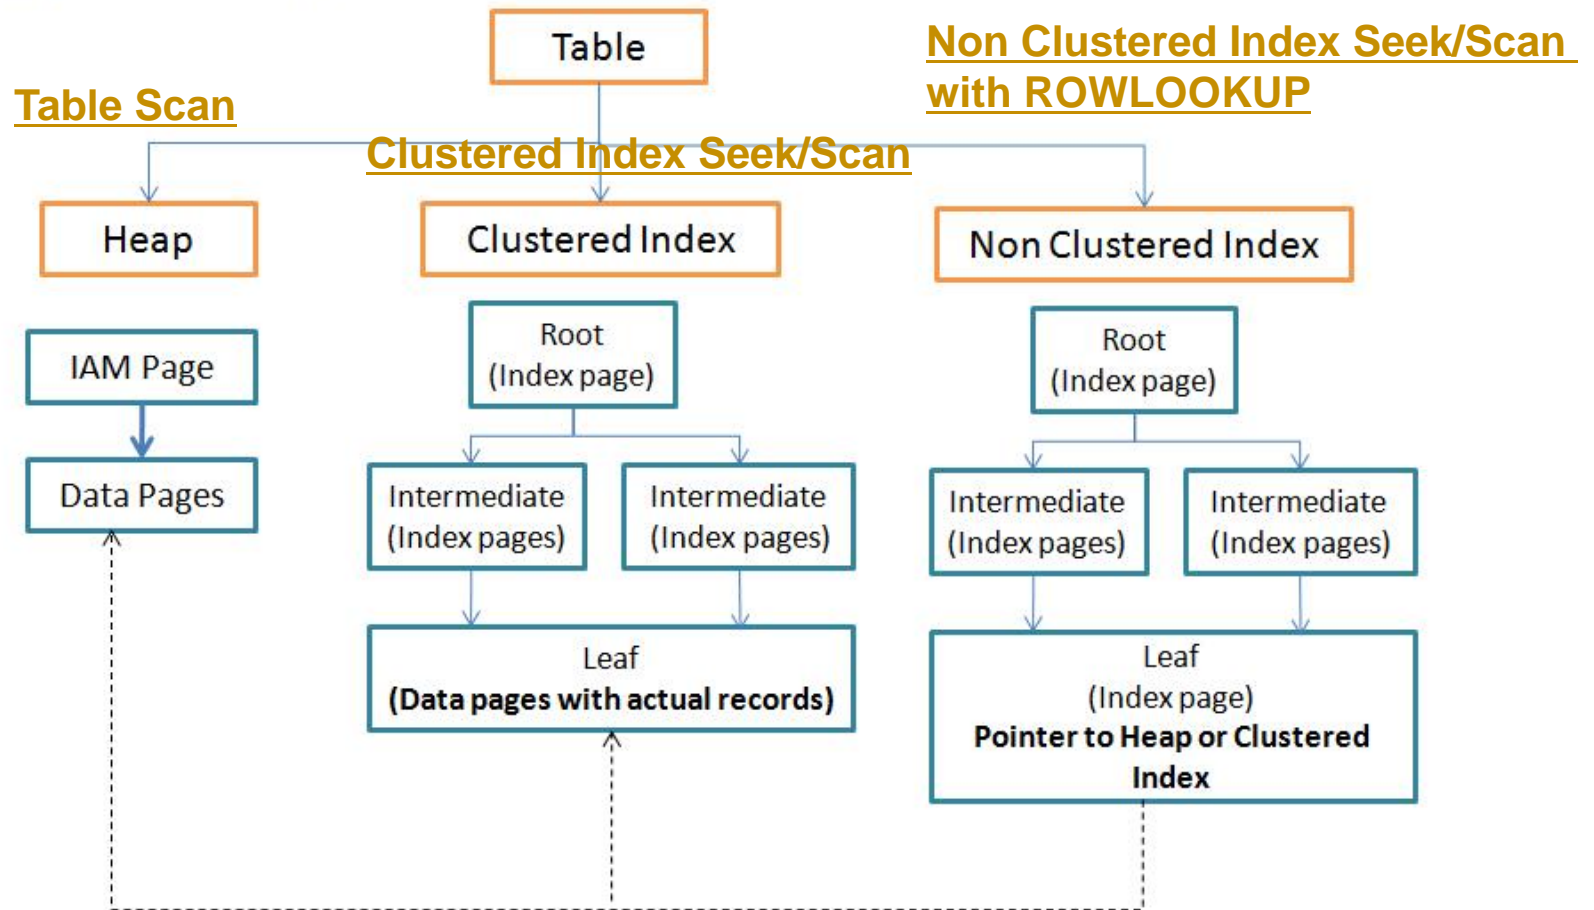
Reading records using Non Clustered Index

Reffering Clustered Index with Non Clustered Index



Sql Server Storage Internals –Index Structure

Pages in different types and levels of index





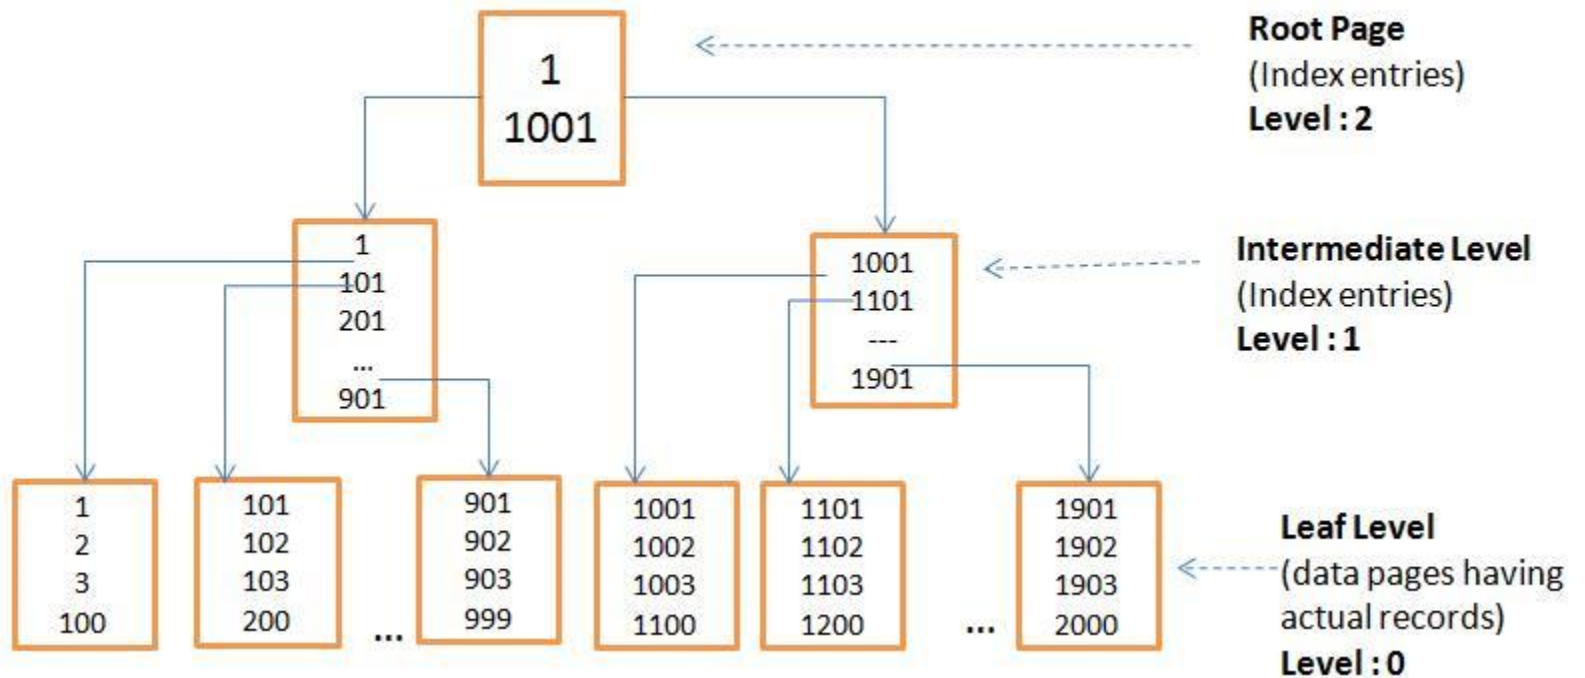
Sql Server Storage Internals – Indexes

Common Properties of Clustered and Non Clustered Indexes:

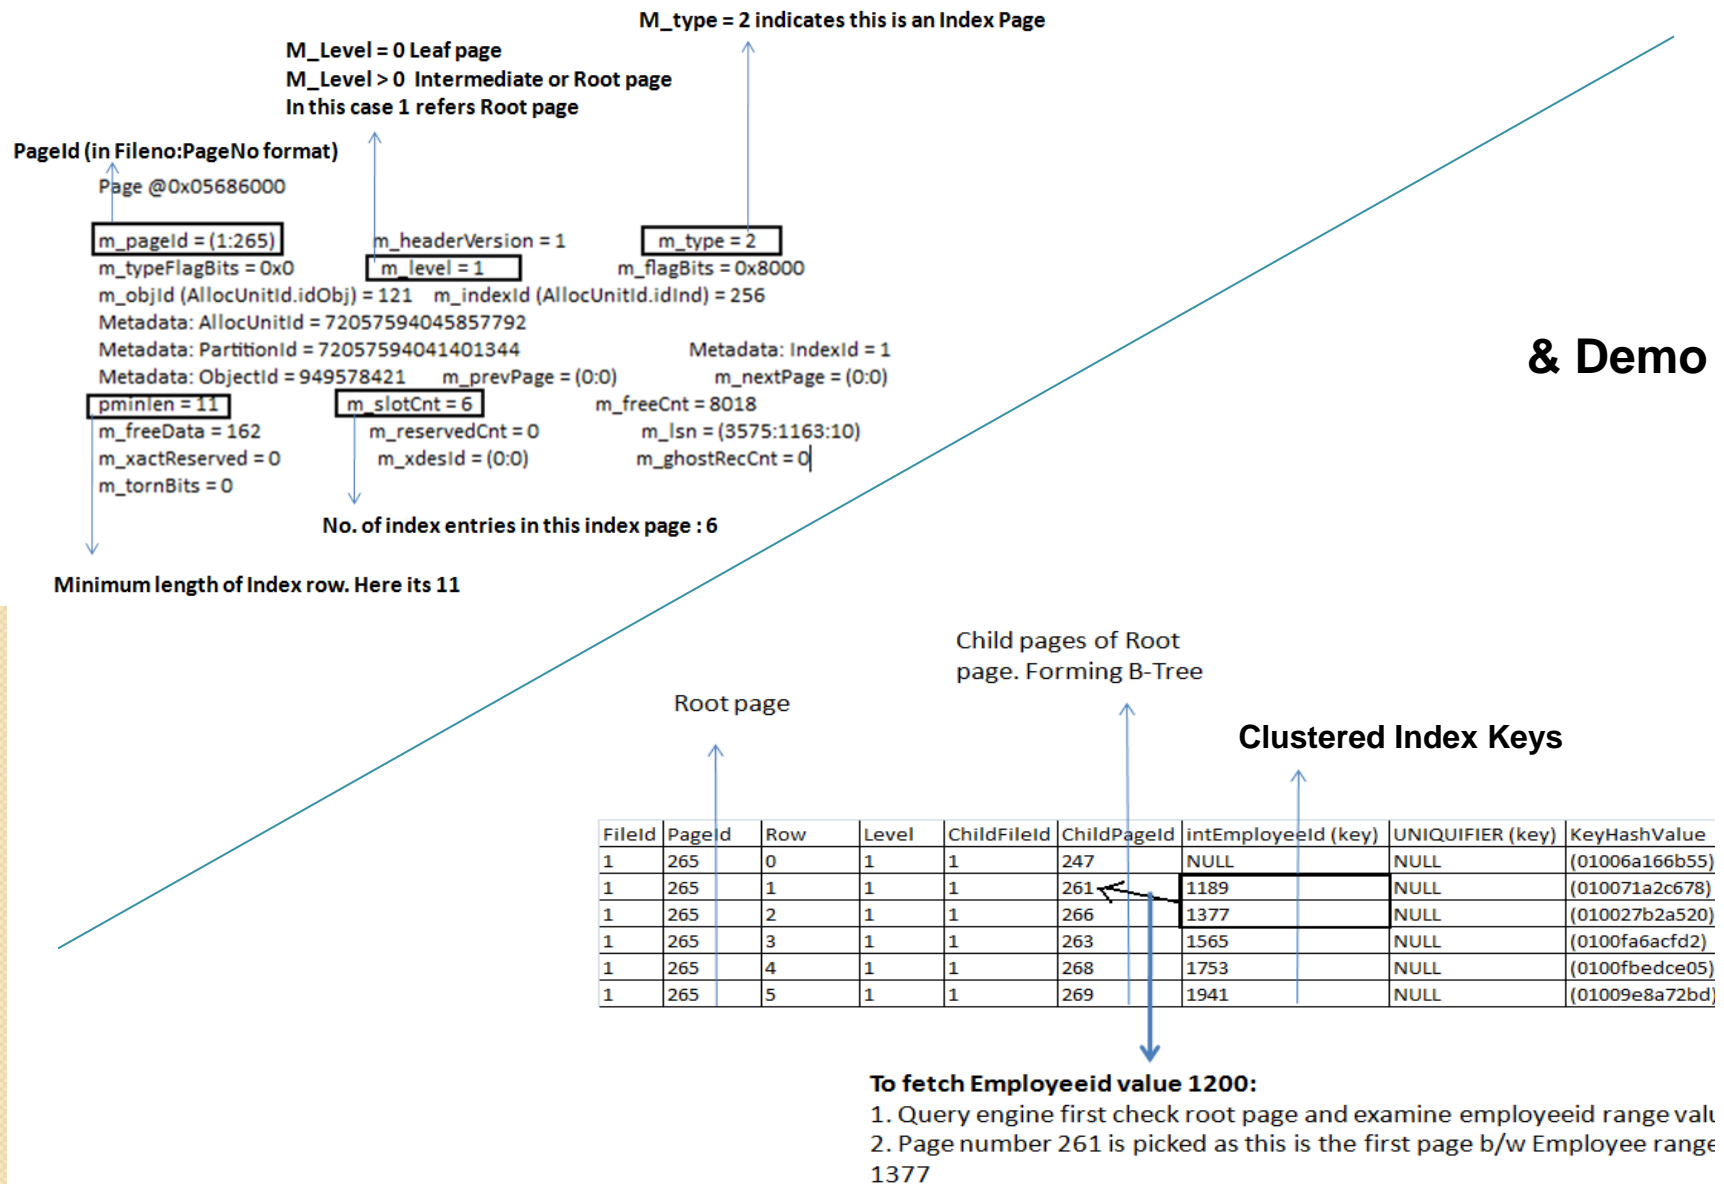
1. Indexes are organized as B-trees.
2. Each page in an index B-tree is called an index node.
3. The top node of the B-tree is called the root node.
4. The bottom level of nodes in the index is called the leaf nodes.
5. Any index levels between the root and the leaf nodes are collectively known intermediate levels.
6. Each index row contains a key value and a pointer to either an intermediate page in the B-tree, or a row in the leaf level of the index.
7. The pages in each level of the index are linked in a doubly-linked list.

Sql Server Storage Internals – Clustered Indexes

Clustered Index Architecture:

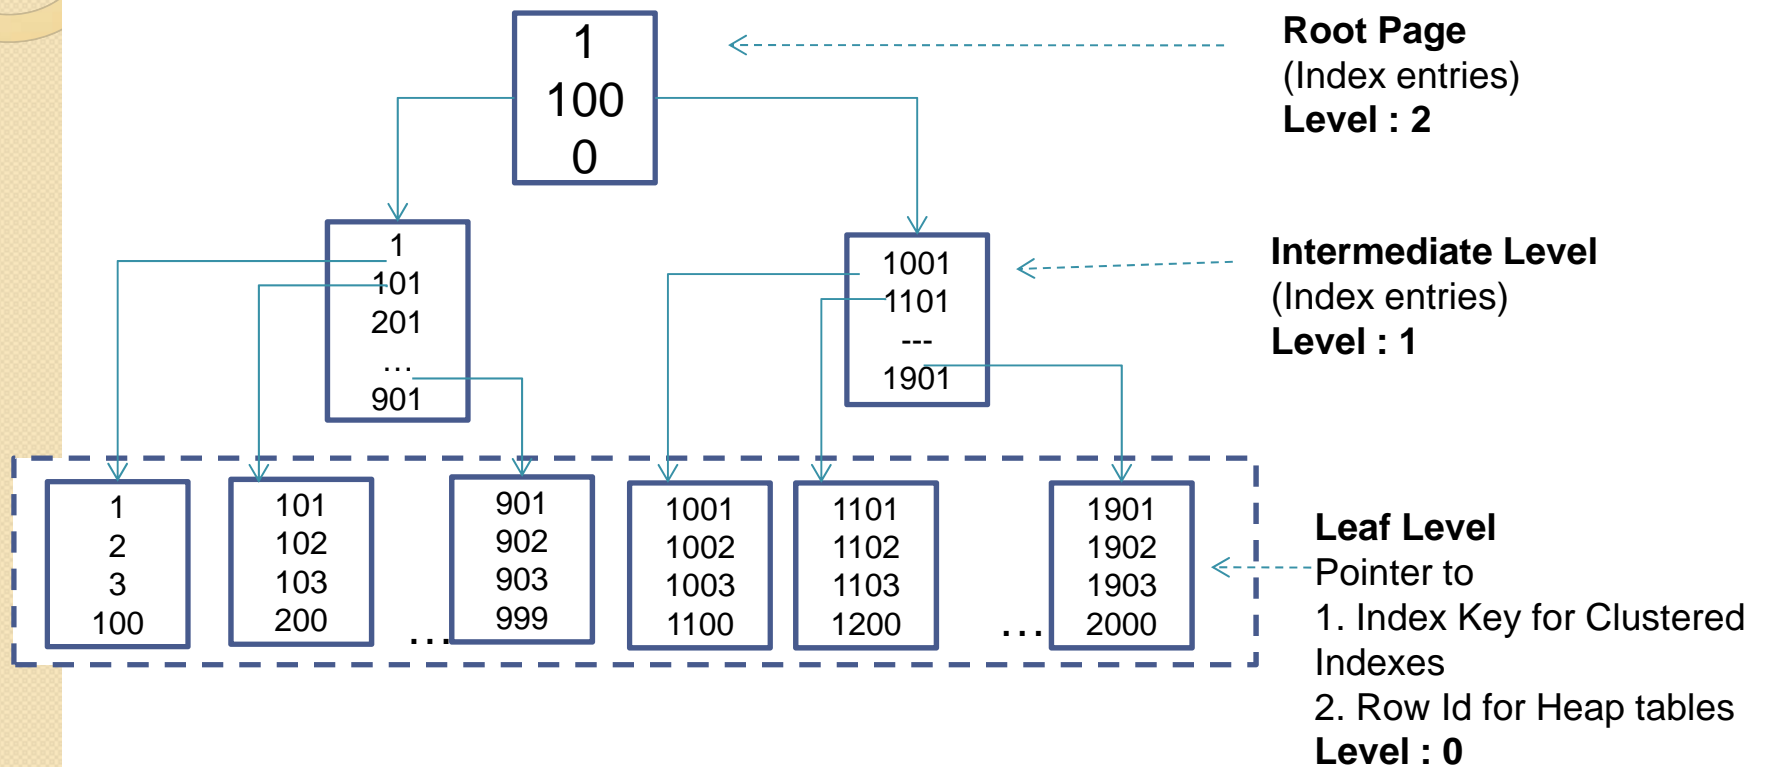


Sql Server Storage Internals – Reading Clustered Index Root Page



Sql Server Storage Internals – Non Clustered Indexes

Non Clustered Index Architecture:



Sql Server Storage Internals – Reading Non Clustered Index Root Page

Non Clustered Index Page of Heap

3 NCI child or leaf pages

NCI root page number 283

NCI is referring Heap table RowId

Field	PageId	Row	Level	ChildField	ChildPageId	strDeptCode (key)	HEAP RID (key)	KeyHashValue
1	283	0	1	1	280	NULL	NULL	(9b009f0ca671)
1	283	1	1	1	282	DEPT48	0x0601000001002300	(b70077a13b0)
1	283	2	1	1	284	DEPT84	0x0601000001009400	(b7009a519fe7)

To fetch Dept value DEPT50:

1. Query engine first check root page of NCI and examine DeptId range values
2. Page number 282 is picked as this is the first page b/w Dept range 48 and 84

Sql Server Storage Internals – Reading Non Clustered Index Leaf Page

Non Clustered Index Page of Heap

Row number

DeptCodes in NCI and Row Ids of Heap to find the row

Field	PageId	Row	Level	strDeptCode (key)	HEAP RID (key)	KeyHashValue
1	282	45	0	DEPT51	0x0601000001009800	(b10069d043f4)
1	282	46	0	DEPT51	0x1201000001008900	(bd00a8cc6b5f)
1	282	47	0	DEPT51	0x1301000001005000	(be0060b23bc3)
1	282	48	0	DEPT51	0x1501000001009400	(c000ada9c2aa)
1	282	49	0	DEPT51	0x150100000100AF00	(c00095467096)
1	282	50	0	DEPT52	0x060100000100A100	(b20016614ac3)
1	282	51	0	DEPT52	0x1201000001007D00	(be00d4d43616)
1	282	52	0	DEPT52	0x120100000100A600	(be0000c8fa74)
1	282	53	0	DEPT52	0x1301000001001F00	(bf006fdd578d)
1	282	54	0	DEPT52	0x1301000001002C00	(bf005fb83c79)
1	282	55	0	DEPT52	0x1401000001004700	(c0002a03f0f5)
1	282	56	0	DEPT52	0x1401000001007600	(c0009804ad33)
1	282	57	0	DEPT52	0x1401000001008F00	(c0006c5ed3f6)
1	282	58	0	DEPT52	0x1501000001001B00	(c100ec11542f)

Index Level 0 refers this page is a leaf page

Sql Server Storage Internals – Reading Non Clustered Index Page

Non Clustered Index Page – Referring Cluster index Key

NCI key (DeptCode) and pointer (Clustered Index key) to Clustered Index

Row numbers of NCI
Leaf page 296

FileId	PageId	Row	Level	strDeptCode (key)	intEmployeeId (key)	UNIQUEIFIER (key)	KeyHashValue
1	296	449	0	DEPT51	1153	0	(2e01e255d486)
1	296	450	0	DEPT51	1326	0	(db00c03a1fcc)
1	296	451	0	DEPT51	1457	0	(5e01018a499f)
1	296	452	0	DEPT51	1901	0	(1a01d2d04579)
1	296	453	0	DEPT51	1928	0	(35017d987199)
1	296	454	0	DEPT52	1162	0	(38015d169077)
1	296	455	0	DEPT52	1314	0	(d0006fa57b8f)
1	296	456	0	DEPT52	1355	0	(f90053b0591b)
1	296	457	0	DEPT52	1408	0	(2e0199699c85)
1	296	458	0	DEPT52	1421	0	(3b0128ad0c7d)
1	296	459	0	DEPT52	1636	0	(120168a77d4a)



Demo